September 1984

Chips Support Two Local Area Networks

BOB DAHLBERG
LAN Component Product Line Manager

Reprinted with permission from May, 1984 issue of Computer Design,

Order Number: 231208-001

SYSTEM DESIGN/ DATA COMMUNICATIONS

CHIPS SUPPORT TWO LOCAL AREA NETWORKS

Data communication ICs permit easy implementation of Ethernet and high level data-link control networks.

by Bob Dahlberg

The main rationale for local networks is resource sharing. Today, small, powerful computers using VLSI components sell for less than $2000. Under the circumstances, companies intending to use several such systems are reluctant to equip each one with a disk drive and printer that could more than double the price per station. Rather, they prefer to share disks and printers among several systems in order to spread the cost of peripherals across several users.

By connecting these small computers to a local area network (LAN), resource sharing with little degradation in overall system performance becomes practical. However, if the network interface costs $1000 or more per computer, the economic advantage of resource sharing wanes. Thus, network interface cost is a primary criterion in selection, particularly for low cost computers.

Access methodologies represent another important factor in network selection. And, although an equal access, first-come, first-served method might be appropriate for an office system environment, it could be the curse of a process control system. In the latter case, a priority-based (or controlled) access method might be the only realistic choice.

Bob Dahlberg is a product manager responsible for local area network components at Intel Corp, 3065 Bowers Ave, Santa Clara, CA 95051. He holds a BS in electrical engineering and computer science from the University of California, Berkeley, and an MBA from the University of Chicago.

COMPUTER DESIGN/May 1984

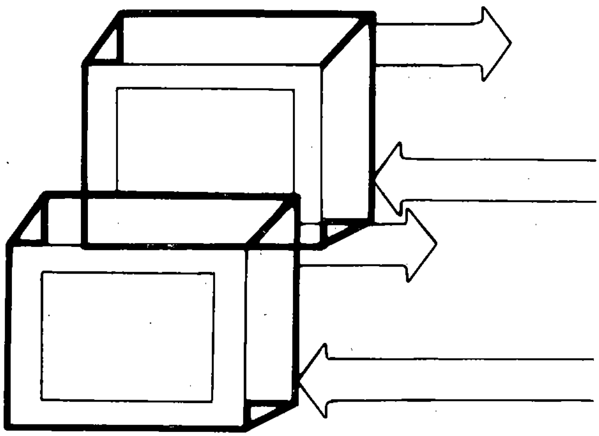

All else being equal, networks supported by available LSI and VLSI components exhibit cost and development speed advantages over board-based LANs. Now, available chips support both priority-based and equal access schemes. One such network is based on the IEEE 802.3 specification, while another uses a variety of physical interface schemes overlaid by high level data-link control/synchronous data-link control (HDLC/SDLC) protocols.

Costly copper

In short distance networks, one can choose a serial, two-wire scheme or a parallel, multiwire interface. Parallel bus structures are implicitly faster than serial structures but tend to be more expensive and less reliable. The amount and cost of the copper

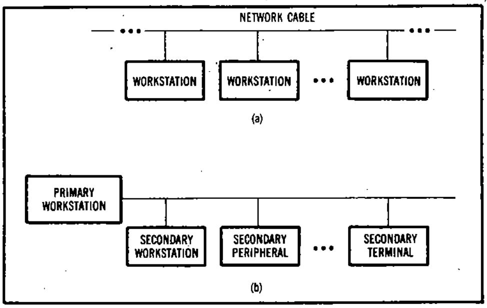

Fig 1 A multidrop configuration is the simplest means of network expansion (a). Additional stations are connected directly to the network cable, but some addressing method must be used to avoid party-line reception by all stations. HDLC/SDLC protocols provide a controlled-access technique where a primary station controls all bus access and determines which secondary stations respond to its commands (b).

wire are much greater, and the number of connections (inversely related to reliability) is also much greater. Thus, the networks described are both serial, two-wire types.

A fundamental assumption in data communications is that noise will corrupt the transmitted data. Error detection schemes can be employed to determine the validity of received data. One common data error detection method applies a numerical algorithm to the message bit pattern and produces a unique sum. This sum is appended to the end of the message and is used by a receiving system as a quick check for the proper bit pattern. Called a cyclic redundancy check (CRC), this process permits a receiving station to discard erroneous data and request retransmission. If the message frames are sequentially numbered, the retransmission request can be made specific to that frame to dispense with the request for a larger group of data. Thus, the process can be made more efficient.

As needs grow, users may want to add more workstations and intelligent peripherals to a network. It would be ideal to attach each station to the network by simply connecting the station directly to the serial network bus cable. This is called a multidrop configuration and it resembles a party line telephone circuit [Fig 1(a)]. As a party line, each station attached to the cable receives all the data transmitted on the cable. In order to route messages to their intended recipients, the messages are logic switched, or specifically addressed, to one or more receiving stations. All others will ignore the data after learning that no match existed between their addresses and those of the data being sent.

Each data packet or frame contains a set of address bits that determines which stations receive the data. In a sense, address bits constitute overhead because they are not part of the information being sent between stations. Any loss in data transfer efficiency, however, is made up by the simplicity of the network expansion interconnect scheme.

The Ethernet specification (a modified version of which was recently accepted as IEEE standard 802.3) describes its physical link characteristics in full detail. Coaxial cable is used as the network cable bus, and each station is connected to that cable via a transceiver and transceiver cable. Minimum distance between station transceivers is 2.5 m, and a network segment can extend to 500 m (and contain up to 200 nodes). Because up to five segments can be joined using active repeaters between each segment, the overall Ethernet network can be 2500 m long and support up to 1000 nodes. Individual nodes can connect to more than one station, and the number of stations connected to an Ethernet network can exceed 1000.

Data is sent at a 10-Mbit/s rate using a self-clocking Manchester encoding format. Only one data packet can be sent at a time using Ethernet, and access is on a first-come, first-served basis. Carrier sense multiple access/collision detection (CSMA/CD) methodology is used. The maximum and minimum distances between transceivers are derived from the CSMA/CD requirements based on interframe-spacing and the collision detection procedures.

A second alternative requires no specific physical link. Speed, distance, and cost parameters dictate actual implementation. The simplest and least expensive method is to drive a twisted-pair cable with off-the-shelf transceiver chips.

Choosing protocols

Both the IEEE 802.3/CSMA/CD and the HDLC/SDLC protocols provide logic-switched messaging and frame-by-frame error detection. IEEE 802.3/Ethernet treats each station equally and does not permit priority network access, whereas HDLC/SDLC enforces a primary/secondary hierarchy [Fig 1(b)]. A primary station controls the overall network by issuing commands to the secondary stations. Secondary stations comply with the primary station's commands and access the bus for retransmitting data only in response to those commands. Unlike Ethernet, which is based on probabilistic network access, HDLC/SDLC provides deterministic (or controlled) access.

SDLC is an IBM standard communication protocol and a subset of HDLC, a standard communication link control established by the International Standards Organization (ISO). HDLC and its subset are data-transparent protocols, which means the arbitrary data streams can be sent without concern that some of the data might be mistaken for control characters. Thus, unlike the Bisync protocol and its controller, an HDLC/SDLC controller need not detect special characters except for the unique opening/closing flag bytes. Moreover, unlike an

OPENING FLAG (F)	ADDRESS FIELD (A)	CONTROL FIELD (C)	INFORMATION FIELD (I)	FRAME CHECK SEQUENCE (FCS)	CLOSING FLAG (F)
01111110	8 BITS	8 BITS	VARIABLE LENGTH (ONLY IN INFORMATION FRAMES)	16 BITS	01111110

Fig 2 The prescribed format for HDLC/SDLC frames consists of four basic fields bounded by opening and closing flags. This avoids the need for start/stop bits often used in asynchronous protocols.

asynchronous protocol and its controller, the HDLC/SDLC need not provide start and stop bits.

Both HDLC/SDLC and Ethernet protocols specify particular message formats (or frames). The HDLC/SDLC protocol consists of five basic fields—flag, address, control, data, and error detection. Each frame is enclosed by an opening and closing flag. Both the opening and closing flags form a similar bit sequence—01111110—that is an individual character in SDLC/HDLC. Inserting a 0 in the information data flow whenever a sequence of five 1s occurs achieves flag character individuality in SDLC/HDLC. These inserted 0 bits are automatically stripped out upon reception. For SDLC, the address field is 8 bits wide, but can be 2 (or more) bytes long in HDLC. Similarly, the control field in SDLC is 8 bits wide, but can also be longer in HDLC. The SDLC data or information field can contain any number of bytes. However, the same is true for HDLC in certain instances where the data field must end on an 8-bit boundary. Finally, the frame check sequence field contains the 16-bit CRC result for all of the bits between flags (Fig 2).

Three types of frames are used in HDLC and SDLC. A nonsequenced frame establishes initialization and control of the secondary stations. A supervisory frame handles control, and an information frame is used for data transfers.

The SDLC protocol appears in low cost asynchronous modems using nonreturn to zero inverted (NRZI) coding and decoding. NRZI coding is used at the transmitter to enable clock recovery from data at the receiver terminal. Clock recovery is accomplished using a digital phase locked loop technique. NRZI coding specifies that the signal condition does not change for transmitting a 1, but changes state whenever a 0 is transmitted. Hence, NRZI coding ensures that an active data line will have a transition at least every 6 bit times (by virtue of the 0-bit insertion requirement). Both 0-bit insertion and NRZI coding/decoding maintain the data transparency characteristics of the HDLC protocol and its SDLC subset.

Like HDLC/SDLC, Ethernet specifies a frame format (Fig 3). It contains a destination field, source field, frame type field, data field, and a frame check sequence. The destination and source fields both contain 6 octets (8 serial bits), for a total of 48 bits. The type field contains 2 octets. The data field can have as few as 46 octets or as many as 1,500. Finally, the frame check sequence consists of 4 octets, allowing a 32-bit CRC code to be calculated and appended to the rest of the frame. The first transmitted Ethernet frame is preceded by a 64-bit preamble, made up of seven groups of 10101010 followed by an eighth group of 10101011. The next bit that follows is the first bit of the first destination octet.

In the CSMA/CD scheme, a "collision" occurs when two stations attempt to gain access to the bus at the same time. Thus, it is important that all stations on the network are notified of the collision. This way, any transmitted data can be flagged as invalid. To solve this problem IEEE 802.3/Ethernet specifies that, after collision detection, transmitting stations send a jam signal to ensure that stations on the network recognize the collision. At the end of the jam interval, each station delays bus access according to an individually calculated random backoff time interval. Should a collision occur again when bus access attempts are renewed, the next backoff interval increases in length. Up to 16 repeated attempts can occur before a system fault is automatically assumed. Thus, even during periods of high bus demand, ample bandwidth should be available and delays relatively short.

It's in the chips

Any of the working LANs can be implemented using various components. If there is enough time and a large budget, custom VLSI chips can be

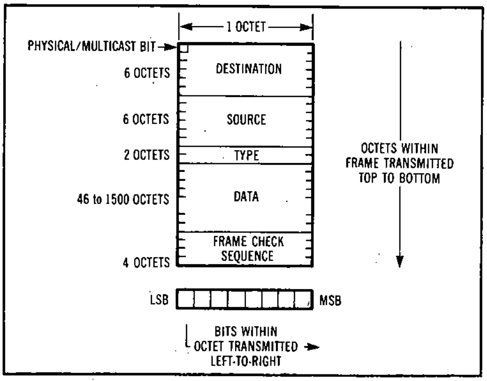

Fig 3 Each Ethernet frame consists of five fields. Destination and source fields indicate where the message is going and from which station it originated. The data field can contain as few as 46 bytes of data and as many as 1500.

developed and an elegant solution forged. Most engineers, however, have neither luxury. For this reason, the two networks selected are supported by off-the-shelf VLSI components.

Intel's 8273 and 8274 data communication controller ICs offer HDLC/SDLC capabilities. Teamed with a microprocessor and some random logic ICs, a capable network data-link controller could be built. The 8051 single-chip microcontroller has become a popular component for many terminal applications because of its high performance 8-bit CPU, large internal program and data memory capacity, plus onchip counter timers and interrupt controllers. In addition, Intel has combined an intelligent HDLC/SDLC controller and 8051 core processor onto a single chip, the 8044. The resulting single-chip microcontroller with onchip serial communication controller allows low cost network terminal and peripheral design.

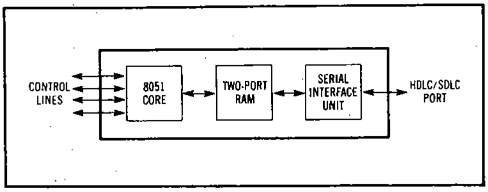

Fig 4 The 8044 combines an 8051 CPU, program and data memory, plus HDLC/SDLC controller on a single chip to build a simple, low cost network station or peripheral.

Each station would contain an 8044 (with its programmable I/O ports to provide local control) and serial HDLC/SDLC interface. Thus, to manage the network interface, 8044-based stations would be capable of acting as a secondary station within an HDLC/SDLC network (Fig 4). Since data transfer speed and electrical characteristics are not specified for these protocols, the designer has a wide choice in tailoring the physical link to the application. The single VLSI device provides local intelligence and network management, thus permitting low cost network development.

Various Ethernet controllers have been announced, with several already sampled and available. Among these is the 82586 general-purpose CSMA/CD controller. It is designed to come up in the Ethernet mode on power up, but can be programmed for other parameters as well. A companion chip (the 82501) provides the Manchester encoding/decoding function between the 82586 and a transceiver.

This chip pair operates in conjunction with the iAPX 86 microprocessor family, and is most cost-effectively used with the 80186 microprocessor. The 80186 and 82586 have identical bus interface and control signal requirements. Hence, they can be linked without adding random logic ICs. Essentially, these three ICs—the 80186, 82586, and 82501—provide the basis for an Ethernet interface. Therefore, only some buffer memory and bus interface chips are additionally required (Fig 5).

A subsystem built using these components provides an intelligent Ethernet interface that can continuously operate at the full 10-Mbit/s network speed. Moreover, these components can implement a complete computer and communication system. It is therefore possible to create an appropriate and usable Ethernet workstation out of these few VLSI components.

Different strokes

The HDLC/SDLC-based network is intended for non-Ethernet applications. HDLC/SDLC has become an accepted standard supported by a variety of hardware and software products. There is no specified standard for physical link implementation or for the software layers beyond the data-link level. Therefore, networks based on these protocols are usually "closed." That is, the vendor provides all the pieces to the network. Vendors, of course, are familiar with their own network architecture and are free to provide compatible systems. But such networks do not encourage others to develop compatible systems unless the vendor's market share is large enough and vulnerable enough to attract competition. The IBM SNA is an example.

HDLC/SDLC-based LANs are suitable for system clusters where distances are less than those of Ethernet, and where priority access is important. Networks within a box (eg, a copier), and networks on table-tops (eg, an instrumentation cluster), are examples. Although there is a parallel bus interface standard (IEEE 488), an instrumentation manufacturer may want to provide for longer distances using two-wire cables and simpler protocols.

An HDLC/SDLC LAN cluster could also be used for process control applications and data acquisition systems. An example is Intel's recent distributed control module products for the factory. Again, a priority bus access capability would be important in these applications. Office system applications where Ethernet offers too much performance at too high a cost (eg, an electronic typewriter networked to a file server) might use this network as well.

The concept of open-system compatibility comes from the ISO's Open System Interconnection (OSI) model. This provides a seven-layer model in which each layer is characterized by a unique set of functions and a specific interface to adjacent layers. The goal is to eventually arrive at a set of standards that would permit systems from several vendors to communicate with one another through common physical, data-link, and software layer protocols.

Xerox Corp developed Ethernet as a local network for its systems, but the company later joined with Digital Equipment Corp and Intel to develop a set of specifications for Ethernet that would allow it to

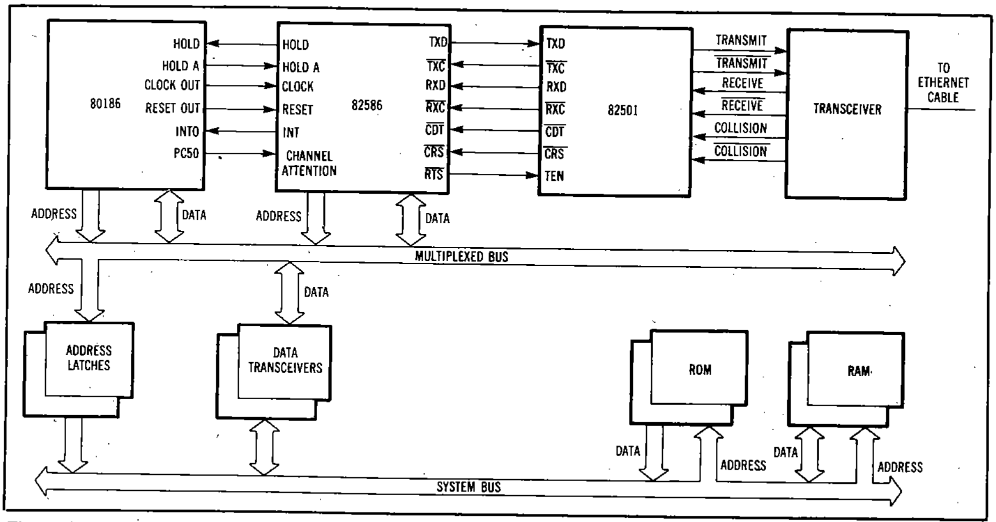

Fig 5 A combination of 80186, 82586, and 82501 chips completes the logic needed for a fully functional Ethernet interface. Data bus interface chips and some memory complete the Ethernet subsystem.

map into the first two layers of the OSI model—physical and data-link. The IEEE adopted its 802.3 specification as a result of these efforts. Efforts to develop standards for the other layers continue. An example is the ISO transport layer protocol, 8073, which provides "return receipt" quality communication services.

Today, Ethernet supports OSI physical and data packet level protocols. It is an emerging technology that is still closer to the top than to the bottom of the learning (and pricing) curve. Nevertheless, many vendors support Ethernet and will no doubt manufacture products equipped to swap data with other Ethernet systems.

Open and closed

Office automation constitutes the biggest apparent application area for Ethernet. The office has traditionally been a multivendor site in which the computer, copier, and printer are likely to come from different vendors. An open system appeals to users seeking vendor independence.

When the LAN concept was first proposed, it was described as an all-encompassing network, connecting all the intelligent subsystems throughout a facility. In fact, that is not the way local network installations have progressed. Instead, clusters of user stations (typically 10 or so) are cropping up in various places within a facility. Most analysts expect local networking to occur in tiers. The cluster tier provides the lowest cost per connection. An example is a 1-Mbit/s CSMA/CD LAN used for personal computers. Clusters would be interconnected through a longer and faster data highway (called a LAN backbone) such as Ethernet.

Will closed and open networks be able to cooperate and coexist? Quite simply, they have to. Economics will determine the network types used for connecting the systems within a cluster, and standardization will drive the methods by which clusters are ultimately joined.

Closed systems, such as microcontrollers connecting the HDLC/SDLC-based network, represent the least expensive and most flexible LAN configuration. Open systems, because of the push for standardization and subsequently larger user base, are more likely to benefit from future cost reduction through multiple-sourced VLSI components than closed systems. Similarly, open systems probably attract more third-party suppliers and enjoy greater variety and lower cost software.

Gateways will join closed and open systems. These hardware/software intermediaries will pave the way for data transfer between formerly incompatible networks. By such means, a closed engineering workstation network will gain access to information stored in the corporate data base and be available on the Ethernet data highway.